

033048-025

**APPLICATION FOR UNITED STATES  
LETTERS PATENT**

**by**

**RAYMOND SUORSA**

**for**

**AUTOMATED PROVISIONING FRAMEWORK  
FOR INTERNET SITE SERVERS**

Burns, Doane, Swecker & Mathis, LLP  
Post Office Box 1404  
Alexandria, Virginia 22313-1404  
(703) 836-6620

Attorney Docket No. 033048-025

**AUTOMATED PROVISIONING FRAMEWORK  
FOR INTERNET SITE SERVERS**

Field of the Invention

The present invention is directed to the provisioning of servers and other computing devices that provide support for sites that are hosted on the Internet, intranets, and other communication networks, and more particularly to a framework that facilitates the automated provisioning of such devices during operations such as the initial deployment of a site, rescaling of the site and/or disaster recovery.

Background of the Invention

The growing popularity and increasing accessibility of the Internet has resulted in its becoming a major source of information, as well as a vehicle for inter-party transactions, in a variety of environments. For instance, a number of different types of entities, such as government agencies, school systems and organized groups, host Internet and/or intranet web sites that provide informational content about themselves and topics related to their interests. Similarly, commercial enterprises employ web sites to disseminate information about their products or services, as well as conduct commercial transactions, such as the buying and selling of goods. To support these activities, each web site requires an infrastructure at one or more centralized locations that are connected to a communications network, such as the Internet. Basically, this infrastructure stores the informational content that is associated with a particular site, and responds to requests from end users at remote locations by transmitting specific portions of this content to the end users. The infrastructure may be responsible for conducting other types of transactions appropriate to the site as well, such as processing orders for merchandise that are submitted by the end users. A significant component of this infrastructure is a web server, namely a computer having

007E01"05E66960

software which enables it to receive user requests for information, retrieve that information from the appropriate sources, and provide it to the requestor. Web sites which provide more complex services, such as online ordering, may also include application servers to support these additional functions.

5           In the case of relatively small entity, the infrastructure to support its web site may be as simple as a single server, or even a portion of a server. Conversely, a large, popular web site that contains a multitude of content and/or that is accessed quite frequently may require numerous web servers to provide the necessary support. Similarly, web sites for commercial entities, via which  
10           transactional operations are conducted, may employ multiple application servers to support transactions with a large number of customers at one time. In addition to servers, the infrastructure for a web site typically includes other types of computing devices such as routers, firewalls, load balancers and switches, to provide connectivity, security and efficient operation.

15           The present invention is particularly directed to the manner in which servers, and other devices necessary to support a web site, are provisioned with the appropriate software necessary for the site. Provisioning includes the installation of the software that is executed by the device to perform the functions assigned to it, and the subsequent configuration of that software to optimize its  
20           operation for the given site. Such provisioning initially occurs when the web site is launched, i.e. when one or more servers are connected to an appropriate communications network such as the Internet, and loaded with the programs and data content necessary to provide the services associated with the site. Thereafter, a need for further provisioning may arise, particularly in the case of a successful  
25           web site, when additional servers must be added to support an increasing number of requests from end users. In another instance, the provisioning of the servers and other computing devices may be required as part of a disaster recovery operation, for example a sudden interruption in power, an attack by a hacker, or corruption of stored software and/or data.





00699350 "103100

with it, i.e. the informational content and services associated with a given site will not be the same as those of any other site supported by that data center. These differences may require a combination of hardware and software which is unlike that of any other site. Similarly, different web site developers may employ  
5 different platforms for the sites, thereby necessitating various combinations of operating systems and application programs on the servers of the respective sites. Furthermore, different types of equipment may be utilized for the sites, thereby adding to the complexity of the provisioning process. In some cases, the same site may require a variety of different hardware devices, operating systems and  
10 application programs to handle all of the different services provided by that site. For an entity that is responsible for managing the varied infrastructure of these sites, such as a data center operator or a third-party infrastructure utility provider, the known approaches to automated software deployment are not adapted to meet the high degree of customization that prevails in these types of situations. Rather,  
15 because of the flexibility that is required to accommodate a different configuration of hardware and/or software for each site, manual provisioning is still being practiced to a large extent, with all of its attendant disadvantages.

It is desirable, therefore, to provide a framework for the automated provisioning of servers and other devices that support various types of network-  
20 based services, such as the hosting of an Internet or intranet web site. Such a framework should exhibit sufficient flexibility to accommodate the differing needs of the hosts of such services, while maintaining repeatability, and hence reliability, in the provisioning process.

#### Summary of the Invention

25 In accordance with the present invention, the foregoing objectives are achieved by means of a framework in which an automated provisioning system communicates with agents that are resident on each device that is to be provisioned, such as servers, routers, and other computing devices. The agents

have access to the configuration of the device at the highest level of authority, so that they are able to fully manipulate all of the software on the device. To minimize security risks in light of the authority level of the agent, communications between the agents and the provisioning system are authenticated, encrypted and carried out in a point-to-point manner.

The provisioning system includes a central file system which contains all of the software components that need to be installed on the devices. In one embodiment of the invention, these components are classified into different categories, or roles, that relate to the frequency with which they are likely to be updated and/or the set of personnel who are permitted to have access to them. Thus, for example, static data content of a site, which may be updated on a daily or weekly basis, is maintained in a separate role from application programs and operating system software. Each device has a set of roles assigned to it, which determines the function and operation of that device.

The system also includes a central database that contains all information that is relevant to the provisioning of the devices. This information includes the hardware configuration of the devices, the software components that make up the various roles assigned to a device, the configuration settings for those components, and logical information such as IP addresses and the like. Whenever a device is to be automatically provisioned, the relevant information is retrieved from the central database, and used to construct a set of commands that are sent to the remote agents for installing the appropriate software components and otherwise configuring the devices.

The information stored in the central database comprises a model of the individual devices, as well as the interconnections of those devices. Whenever a change is to be made to a device, the change is first recorded in the stored model, rather than being directly implemented on the device itself. These changes are carried out by means of a user interface that enables an operator to test them on the model. Once the changes have been verified to be appropriate, they are then

provided from the database to the device, through the agents. By having the configurations of the devices be controlled from the database, rather than directly by operators, repeatability of results is assured for all devices of the same type.

5 All communications between the central database and the remote agents are preferably carried out by means of a central gateway within the provisioning system. This gateway converts provisioning policies from the user interface and database information into the primitives of messages that are sent to the remote agents. As a result, the agents themselves can be relatively light weight in structure, and need not possess a significant amount of internal functionality to  
10 perform the tasks associated with provisioning the devices.

These and other features of the invention are explained in greater detail hereinafter with reference to an exemplary embodiment of the invention illustrated in the accompanying drawings.

#### Brief Description of the Drawings

- 15 Figure 1 is a block diagram of the basic logical tiers of a web site;  
Figures 2a and 2b are more detailed diagrams of the devices in an exemplary web site;  
Figure 3 is a block diagram of one embodiment of the hardware configuration for a web site in a data center;  
20 Figure 4 is a more detailed block diagram of an exemplary configuration for a web site host compartment in a data center;  
Figure 5 is a time line illustrating the life cycle of a typical web site server;  
Figure 6 is a general block diagram of a data center in which the present invention can be implemented;  
25 Figure 7 is a block diagram of a provisioning framework in accordance with the principles of the invention;  
Figure 8 is a block diagram of the roles for server software;  
Figure 9 is a diagram of the hierarchy of components in a role;



Figure 10 is a timing diagram that illustrates the communication between the gateway and an agent; and

Figure 11 is a block diagram of the components of the agent.

#### Detailed Description

5 To facilitate an understanding of the principles of the present invention, it is described hereinafter with reference to its application in the provisioning of devices that support web site operations, such as servers, load balancers, firewalls, and the like. Further in this regard, such description is provided in the context of a data center, which typically accommodates the infrastructure to support a large  
10 number of different web sites, each of which may have a different configuration for its infrastructure. It will be appreciated, however, that the implementation of the invention that is described hereinafter is merely exemplary, and that the invention can find practical application in any environment where the automated provisioning of computer resources is desirable. Thus, for example, the principles  
15 which underlie the invention can be employed to provision computing devices in the networks of an enterprise, or in any other situation in which there are a sufficient number of computing devices to realize the benefits of automated provisioning.

Prior to discussing the specific features of an exemplary embodiment of the  
20 invention, a general overview of the infrastructure for hosting a web site will first be provided. Fundamentally, a web site can be viewed as consisting of three functional tiers. Referring to Figure 1, one tier comprises a web server tier 10. The web server is the combination of hardware and software which enables  
browsers at end-user locations to communicate with the web site. It performs the  
25 task of receiving requests from end users who have connected to the web site, such as HTTP requests and FTP requests, and delivering static or dynamic pages of content in response to these requests. It also handles secure communications through a Secure Socket Layer (SSL), and the generation of cookies that are

downloaded to browsers. Typically, since these types of operations do not require a significant amount of processing power, the web server can operate at relatively high volume rates. The throughput capacity of this tier is usually determined by the amount of server memory and disk storage which is dedicated to these operations.

Another tier of the web site comprises an application server tier 12. This component performs dynamic transactions that are much more computationally intensive, such as order processing, credit card verification, etc. Typically, the application server implements the development environment that defines the business logic and presentation layer associated with a given site, i.e. its functionality as well as its "look and feel". The performance of this tier is normally determined by the amount of CPU processing power that is dedicated to it. Separation of the web servers and the application servers into different tiers ensures reliability and scalability.

The third tier of the site comprises a database tier 14. This tier stores information relevant to the operation of the site, such as customer demographic and account information, available stock items, pricing, and the like. Preferably, it is implemented with a relational database architecture, to permit the data to be manipulated in a tabular form. Connection pooling to the database can be performed by the application servers, to minimize redundant calls and thereby preserve processing power.

While the fundamental architecture of a web site can be viewed as comprising these three tiers, in an actual implementation the structure of the web site can be significantly more complex. Depending upon the size and requirements of the site, in some cases the database tier can be combined into the application server tier. Even more likely, however, is an architecture in which one or more tiers is divided into several layers. This occurrence is particularly true for the application server tier, because it implements the business logic of a site. Depending upon the types of transactions to be performed by the site, the

application server tier may require a number of different types of specialized application servers that are interconnected in various ways. One example of such is depicted in Figure 2a. In this situation, the site includes a number of web servers 11a, 11b, ...11n. Each of these web servers may have the same software and same configuration parameters. The site also includes a number of application servers 13a, 13b, ...13n. In this case, however, not all of the application servers are the same. For instance, server 13a communicates with a first type of database server 15a, whereas servers 13b and 13n communicate with another application server 13d at a different level, which may be a highly specialized server. This server may communicate with a second type of database server 15b to carry out the specialized services that it provides. In addition, the server 13n may communicate with a directory server 15c.

If the performance of the server 13d begins to degrade due to increased traffic at the web site, it may be necessary to add another server 13d', to provide additional CPU capacity, as depicted in Figure 2b. However, because of the architecture of the site, the automated provisioning task becomes more complex, since the application server 13d is different from the other application servers 13a, 13b, etc., in both its configuration and its connection to other devices. Hence, not all of the application servers can be treated in the same manner. Furthermore, since the business logic of a given site is likely to be different from that of other sites, the configuration parameters that are employed for the site of Figure 2a may not be appropriate for the devices of any other site, which increases the complexity of the provisioning process even more.

In many instances, the infrastructure for supporting a web site is housed in a data center, which comprises one or more buildings that are filled with hundreds or thousands of servers and associated equipment, for hosting a large number of different web sites. Typically, each floor of the data center contains numerous rows of racks, each of which accommodate a number of servers. In one configuration, each web site may be assigned a portion of a server, or portions of

several servers, depending upon its requirements. This approach is typically employed by Internet service providers (ISPs), and is referred to as a "multi-tenancy" configuration, wherein multiple sites may be resident on a given server.

5 In an alternate configuration, each site is allocated a discrete compartment within the data center, with the servers and other computing devices within that compartment being dedicated to hosting the services of the given site. Figure 3 is a block diagram illustrating this latter configuration. This figures illustrates three exemplary web site compartments, each of which accommodates the equipment for hosting a web site. Thus, in the illustrated embodiment, each compartment

10 includes one or more web servers 10a, 10b, one or more application servers 12a, 12b, and a database server 14a, to provide the three functional tiers. In addition, the components of the web site infrastructure may include a firewall 16 to provide security against attacks on the site, a load balancer 18 for efficient utilization of the web servers and the application servers, and a switch 20 for directing incoming

15 data packets to the appropriate servers. These devices in the web site compartment can be securely connected to the host entity's computer system via a virtual private network 22. To avoid a single point of failure in the web site, additional redundant components are included, and like components are cross-connected with one another. This feature of redundancy and cross-connection adds another layer

20 of complexity to the automated provisioning process, particularly as the web site grows so that the number of devices and their cross-connections increase and become more complicated to manage.

The physical storage devices for storing the data of a web site can also be located in the compartment, and be dedicated to that site. In some cases, however,

25 for purposes of efficiency and scalability, it may be preferable to share the data storage requirements of multiple compartments among one another. For this purpose, a high capacity storage device 24 can be provided external to the individual compartments. When such a configuration is employed, the storage device 24 must be capable of reliably segregating the data associated with one

compartment from the data associated with another compartment, so that the different hosts of the web sites cannot obtain access to each others' data.

Examples of storage devices which meet these requirements are those provided by EMC Corporation of Hopkinton, Massachusetts. For additional discussion of the manner in which devices of this type can be incorporated into an infrastructure such as that depicted in Figure 3, reference is made to co-pending, commonly assigned Application No. \_\_\_\_\_ [Attorney Docket No. 033048-008], filed on an even date herewith, the disclosure of which is incorporated herein by reference.

In a particularly preferred embodiment, each web site compartment is comprised of at least three racks 26 within a data center. Referring to Figure 4, the two outer racks 26a and 26c contain the components of the three basic tiers for a web site. Thus, each rack may contain one or more webserver and/or application servers. The center rack 26b contains the devices associated with interfacing the web site server to external networks. Hence, the necessary switches, firewalls and load balancers are contained in this rack, where they can be easily connected to the servers in each of the two adjacent racks.

To provide the services associated with a web site, each of the servers and other devices in a compartment must be configured with the appropriate software, and then regularly maintained to provide updates consistent with changes in the web site. A typical life cycle for a server is depicted in Figure 5. Referring thereto, after a server has been constructed it is typically delivered to a data center, or other site where the web site's infrastructure is housed, with only the computer BIOS (Basic Input/Output System) installed on it. When it is to be put into operation, it is assigned to a designated web site compartment, and then customized for the tasks that are to be performed for that site. At the outset, an appropriate operating system and other general software are loaded onto the server at Step 1. If desired, the operating system and general software can be pre-loaded onto the server, before it is assigned to a specific compartment. One technique for

00669350-103100

preparing servers ahead of time with an operating system and other general software, so that they are ready for assignment to a compartment and immediate loading of site-specific software, is described in co-pending Application No.

\_\_\_\_\_ [Attorney Docket No. 033048-007], filed on an even date herewith,  
5 the disclosure of which is incorporated herein by reference.

The next major step in the customization of the server comprises the loading of the appropriate software applications that will handle the transactions associated with the web site. Examples of such programs include WebLogic application server distributed by Bea Systems, Inc., and Apache Web Server  
10 provided by The Apache Software Foundation. Once these types of programs have been installed at Step 2, they typically must be configured, i.e. various operating parameters must be set to appropriate values, which is depicted as Step 2a in Figure 5. Thereafter, the data content which is specific to the web site is loaded at Step 3, and further configuration may be carried out at Step 3a. Once all of this  
15 software has been appropriately installed and configured, the web site is launched and continues to run at Step 4. Even after the launch of the web site, however, continued maintenance of the server is required, to accommodate changes in the content of the site, upgrades to application software, and the like. These maintenance cycles may include the installation of software components, as  
20 depicted by the loops which return to Steps 1, 2 and 3 in the life cycle illustration of Figure 5. In addition, regular configuration changes may need to be made, to increase the performance of the site, as indicated by the loops which return to Steps 2a and 3a.

At the end of its life cycle, the server may be decommissioned at Step 5.  
25 This could occur as a result of shutting down the web site, performing a major overhaul of the web site infrastructure, and/or upgrading to new equipment. At this point, the software is removed from the server, whereupon it can be assigned to a new compartment, reloaded with new software appropriate to the existing compartment, or simply retired from use.

One feature of the present invention comprises a system for automating the configuration and maintenance of servers during the entirety of their life cycles, as depicted in Figure 5. Further in this regard, an objective of the invention is to provide a framework to deploy and configure software on a large number of servers within one or more data centers, that may be associated with different respective web sites, and therefore have different provisioning requirements.

An overview of one environment in which the present invention operates is depicted in Figure 6. A data center 28 is partitioned into multiple customer compartments 29, each of which may be arranged as shown in Figure 4. Each compartment is connected to a backbone 30 or similar type of common communication line for access by computers which are external to the data center. For instance, if the compartments are associated with Internet web sites, the backbone 30 constitutes the physical communication path via which end users access those sites over the Internet. The backbone may also form the path via which the web site hosts can securely communicate with the devices in their individual compartments, for instance by virtual private networks.

Also located in the data center 28 is a provisioning and management network 31. This network may be located within another compartment in the data center. This network is connected to the computing devices in each of the compartments 29 which are to be managed. In the embodiment of Figure 6, the provisioning network 31 is illustrated as being connected to the compartments 29 by a network which is separate from the backbone 30. In an alternative implementation, the provisioning network can communicate with the compartments over the backbone, using a secure communications protocol.

The provisioning network 31 may be operated by the owner of the data center, or by a third-party infrastructure utility provider. While Figure 6 illustrates all of the compartments being connected to the network 31, this need not be the case. To this end, multiple provisioning networks may be located in the data center, with each one operated by a separate entity to provision and manage





information about each of the devices that are connected to the provisioning network. Hence, depending upon the extent of the provisioning system, the central database might contain information about devices in only a few web site compartments, or an entire data center, or multiple data centers. The information stored in this database comprises all data that is necessary to provision a device. For instance, it can include the hardware configuration of the device, e.g., type of processor, amount of memory, interface cards, and the like, the software components that are installed on the device along with the necessary configuration of each of those components, and logical information regarding the device, such as its IP address, the web site with which it is associated, services that it performs, etc. For a detailed discussion of an exemplary model of such a database for storing all of the relevant information, reference is made to co-pending Application No. \_\_\_\_\_ [Attorney Docket No. 033048-012], filed on an even date herewith, the disclosure of which is incorporated herein by reference. In essence, the information stored in the database constitutes a model for each device that is managed by the provisioning system, as well as the interconnection of those devices.

The second principal function of the provisioning network is implemented by means of a central file system 34, which is accessed via a file server 35. This file system stores the software that is to be installed on any of the devices under the control of the provisioning system. To facilitate the retrieval of a given item of software and forwarding it to a destination device, the software components are preferably stored within the file system as packages. One example of a tool that can be used to create software packages for a Linux operating system is the Red Hat Package Manager (RPM). This tool creates packages in a format that enables the contents of a package, e.g. the files which constitute a given program, to be readily determined. It also includes information that enables the integrity of the package to be readily verified and that facilitates the installation of the package. To support a different operating system, a packaging tool appropriate to that



provisioning network can first be authenticated by means of a signed certificate, to confirm to the recipient that the sender of the message is a trusted entity.

To further enhance the security of the communications between the provisioning network and the agents, the network includes a central gateway 38 for communications. For instance, when the provisioning of a device is to be carried out, the gateway can send a message to the agent 36, informing it of a task to be performed, e.g. a package to be retrieved and installed. Upon receiving this task, the remote agent contacts the central file server 35, with the address of the package that is needed. Before releasing the package to the agent, the file server can check with the central database to determine whether the agent should have access rights to that package. Preferably, the communication between the central file server 35 and the database 32 is also carried out via the gateway 38. By having all communications with the central database pass through the gateway, it becomes possible to store a trust hierarchy 37 that is associated with the gateway and that identifies the level of trust to be accorded each device that sends messages. Thus, for example, since the file server is located within the a secure network 31, it can be identified in the trust hierarchy 37 as having a high level of trust. Conversely, since the agents are located outside of the network, they have a less trusted level, and therefore appropriate security measures are taken. For instance, an agent may have the ability to make reports about the configuration of its device, but it cannot request information from the central database relating to its device. Similarly, an agent installed on a device for one web site cannot have access to packages in the file system 34 that are affiliated with another web site. This is enforced by having the file server 35 contact the communication gateway to confirm that an entity which requests access to a particular package has the right to do so. In addition to levels of trust that are associated with various devices and software components, the trust hierarchy can also identify personnel who are authorized to access and/or manipulate the components of the system.

Another advantage associated with the use of the central gateway for communications between the agents and the database is that it permits the agents to be light weight, i.e. to have a relatively simple command set. More particularly, if the agents were to directly communicate with the central database, they would  
5 need to have functionality analogous to that of a client in a client-server database system. For instance, if the database utilizes a Structured Query Language (SQL) server, the agents would need to be able to construct SQL queries to send to the database. However, by interposing the gateway as a logic layer between the agents and the database, the necessary functionality for communicating with the  
10 database can be contained solely within the gateway, and the communications between the gateway and the agents can be much simpler. For example, messages that are exchanged between the gateway and the agents can be in the form of remote procedure calls that conform to the XML-RPC protocol, or the Simple Object Access Protocol (SOAP). When a message is received at the gateway, it  
15 constructs an appropriate SQL query, to retrieve the appropriate information from the database. This information can then be provided to the agents using a higher level messaging protocol, such as XML-RPC or SOAP. An advantage of using such a protocol is that it enables commands to be sent to the agents from the provisioning network 31, which is not possible with SQL statements.

20 There may be situations in which it is desirable to permit personnel who do not have access to the provisioning system per se to communicate with the agents. For instance, IT personnel at the entity hosting the site may need to perform some types of operations through the agent. In this case, the agent can be given the ability to communicate with a computer 39 external to the network, for instance by  
25 means of a browser on that computer. This external access can also serve as a debugging mechanism. For instance, a new configuration can be set up on a device and then tested in isolation on that device, via the browser, before it is deployed to all of the other devices of that same type. Whenever access to a device is sought by an entity outside of the secure network 28, the agent







However, access to the OS roles may be limited to certain personnel at the data center or other entity which manages the web site infrastructure. The access rights associated with the different roles can be stored in the trust hierarchy 37.

Although the foregoing example has been provided with reference to three  
5 types of roles, it will be appreciated that a greater number of roles can be employed to provide finer gradations between the different types of software on a device. Similarly, it may be preferable to utilize a greater number of roles if more than three different levels of access are set forth in the trust hierarchy for the software components.

10 When provisioning is to be carried out on a device, the commands to perform this operation are provided to the agent 36 for the device by means of a command queue. Each queue comprises a set of commands that are to be run by the agent 36 in a specific order. The commands may be individually designated via the user interface 40, or be a predefined script that is stored in the database 32  
15 and called up via the user interface. The command queue is stored in the database 32 to provide persistence, so that in the event the gateway should experience a failure while a series of commands is being carried out, the queue will still exist when the gateway is restored to an operational state. While a command queue is being executed, the gateway keeps track of its state in the database, i.e. which  
20 command was the last one to be sent to the agent, so that it can easily return to that command if a failure occurs.

The commands are executed via interaction between the gateway 38 and the agent. Referring to Figure 10, once a command queue has been created, the execution of the commands begins with a poke message 42 from the gateway to the  
25 agent, informing the agent that there is a command to be run. The agent opens a new connection to the gateway and returns a response 44 through this connection, inquiring about the command. By requiring the agent to open a new connection, i.e. a different secure socket, and request the command from the gateway, the possibility of spoofing is decreased. Specifically, if a spoofer should attempt to



send an agent a rogue command, the agent will respond to the gateway with a request for a command. If there is no legitimate command to be run by that agent, the gateway simply responds with "No Command", and the agent returns to its prior state.

5           Upon receiving the inquiry from the agent in response to a poke message, the gateway retrieves the first command in the queue, and provides it to the agent in a message 46, e.g. get and install a package at a designated address in the file system. The agent runs the command, and then reports back to the gateway with a result 48. If it takes some time to execute the command, the report message may  
10       be by means of a new socket, to prevent an open interface between the agent and the gateway. The report also includes an inquiry as to the next command to be executed. If there is another command in the queue, it is retrieved by the gateway and forwarded to the agent, e.g. configure the files that were just installed with designated parameter values. The process continues in this manner, until the end  
15       of the queue is reached, at which time the gateway responds to the agent's most recent inquiry with a message 50 that there is no command to be executed. At this point, the procedure ends.

          One of the commands 52 that can be sent to the agent is to reboot its device. In response to receipt of this command, the agent sends a result message  
20       54 which informs the gateway that it is rebooting. The gateway does not respond to this message, but places the command queue in a reboot status. Upon rebooting, the agent sends a message 56 to the gateway to inform it that it has just rebooted. In response, the gateway checks the command queue and, if there are  
      commands remaining to be executed, sends the next command 58 in the queue to  
25       the agent.

          The agent 36 can include functionality for determining the hardware and/or software configuration of the device on which it resides. This feature is useful in identifying discrepancies between the data stored in the database 32 regarding the model for the intended configuration of the device, and the actual configuration of



library contains a number of components 70 that relate to the different functions that are performed by the agent, such as load packages, establish a network connection, etc. These components 70 are generic to all operating systems. Plug-in modules 72 which are specific to the particular operating system 62 are  
5 associated with the library components 70. These plug-in modules communicate with the abstraction layer 60 to cause specific actions to be performed by the operating system. In some cases, the plug-in modules may have the capability to communicate directly with the operating system 62, in which case they can bypass the abstraction layer.

10 The foregoing description has been provided in the context of one provisioning network that may be used to control devices at one data center. It will be appreciated that such a network can be a subnetwork in a wide-area network which controls devices at several data centers. In such an embodiment, the communication gateways in each subnetwork can exchange information with  
15 one another regarding the data stored in their respective database systems 32 and/or software packages in their file systems 34. Hence, if an entity has its web site infrastructure apportioned over several data centers, the provisioning operations can be coordinated amongst the various centers.

From the foregoing, therefore, it can be seen that the present invention  
20 provides a framework for the automated provisioning of devices which constitute the infrastructure of a web site, such as servers. Two significant features of this framework are its flexibility and the repeatability of the results that are obtained. The flexibility permits the varied needs of different web sites to be readily accommodated, and thereby avoids the limitation of having to configure the  
25 architecture of every site the same way. The repeatability ensures that every server will have the proper set of software components once it has been provisioned, and thereby be ready to operate immediately. In addition to these features, the automated provisioning that is provided through this system achieves

a significant time savings, enabling the entire process to be accomplished in substantially less time than is required for manual provisioning.

It will be appreciated by those of ordinary skill in the art that the present invention can be embodied in other forms without departing from the spirit or essential characteristics thereof. For instance, while an exemplary embodiment of the invention has been described in the context of provisioning web site servers in a data center, it will be appreciated that the principles underlying the invention can be applied in any environment where computing devices need to be configured and/or updated on a relatively large scale. The foregoing description is therefore considered to be illustrative, and not restrictive. The scope of the invention is indicated by the following claims, and all changes that come within the meaning and range of equivalents are therefore intended to be embraced therein.